

A retrieval strategy for interactive ensemble data assimilation

Ross N. Hoffman*

August 9, 2011

1. Introduction

Is it better to use radiances or retrievals in data assimilation? I started out my career believing in radiance assimilation (Hoffman 1983; Hoffman and Nehrkorn 1989), but now, with radiance assimilation having become mainstream, I'm having second thoughts. Here I outline an approach following Rodgers (2000) that might have some advantages that uses retrievals along with the so-called Jacobian or sensitivity matrix, actual averaging kernel (AK), actual prior, and estimated retrieval noise covariance. (See § 4.) Others have considered using retrievals along with other information provided by the retrieval process (Joiner and da Silva 1998). Eyre and coworkers used 1d-VAR retrievals within 3d-VAR and 4d-VAR assimilation systems (Eyre et al. 1993; Szyniel et al. 2004; Pavelin et al. 2008).

Using retrievals gives us several advantages (Rodgers 2000, Section 8.3). First it reduces complexity in the H -operator, reduces data volume, allows sequential processing within the retrieval algorithm, allows arbitrary cloud clearing methods, and makes the assimilation system more modular. (Sequential processing means, for example, first use one band for temperature retrievals, then use some other band for humidity retrievals, etc.)² (More in § 2.) Using EOFs gives us additional advantages, including reduced data volume. (See § 8.) There is also the potential to reduce vertical interpolation errors that arise due to the different vertical grids used in the forward problem and the meteorological model. (See § 9.) Note that if the retrieval method uses an EOF representation, then formally the error covariance matrix in physical space is not full rank. (But see below in the discussion of “EOF covariances” where we describe how we create the physical space error covariance matrix at AER.) Finally the AK provides the opportunity to remove the influence of the prior, and allows for interactive retrievals (where, for example, the prior is the ensemble mean forecast). (See § 5.)

¹*Contact information:* Dr. Ross N. Hoffman, Atmospheric and Environmental Research, Inc., 131 Hartwell Avenue, Lexington, MA 02421-3126 Email: ross.n.hoffman@aer.com.

²Why do retrieval people like to do this? I think because of uncertainty in the spectroscopy and in estimating errors, and because the optimal coupled problem is so sensitive.

The motivation for this approach is for instruments with thousands of channels. This includes AIRS, IASI and the Aura TES. In these cases, instead of using thousands of channels or having to decide which are the optimal channels to use in the DAS, these issues would be pushed back to the retrieval scheme. But this approach can be applied more generally, and has the advantage of explicitly treating correlations of errors for a single retrieved profile, with only modest impacts on the DAS.

There are three successive developments in what follows. First, in § 4, we show how to use the AK and the retrieval noise covariance to transform the retrieved quantities into observations that are unbiased and have uncorrelated errors, and to eliminate both the smoothing inherent in the retrieval process and the effect of the prior. Second, in § 8, we show how to transform this result into EOF space, when a truncated EOF series has been used in the retrieval process. This provides a degree of data compression and eliminates those transformed variables that have very small information content. In both approaches a vertical interpolation from the dynamical model coordinate to the radiative transfer coordinate is required. In the third development, in § 9, we propose using the EOFs to essentially reduce the error of the vertical interpolation.

2. Radiance assimilation :: pros and cons

The advantages of direct radiance assimilation are that we use the data in its basic state and there are weak or no error correlations. The disadvantages are that there may be huge numbers of channels, and that what I’ve termed geophysical biases can effectively induce observation error correlations. For example, the meteorological model might have a very simple representation of surface emissivity, while the true surface emissivity is governed by geologic parameters that have long length scales. This difference must be considered an observational error in data assimilation systems (DAS).³ A retrieval scheme might employ a slack variable (*e.g.*, cloud fraction) or use information from some channels to model the emissivity in other channels. Further, spectroscopic inconsistencies may introduce correlations in the forward problem, which enter the problem in the same way as observational error correlations. In the end, I expect that whatever can be done in a retrieval scheme can probably be done in DAS if we are clever enough, but at a cost in complexity and maintenance. I believe it is best to avoid the use of radiances in the DAS and instead use retrievals with a strategy along the lines outlined here whenever an intricate retrieval schemes is favored by the remote sensing community to extract the maximum information content from the observations.

3. Notation

As we will quote so many results from Rodgers (2000), we will adopt his notation, defining terms as we go. However, this means that x discussed here, the vector of retrieved quantities⁴, become

³Like “fish” [one fish, two fish . . .] “DAS” is both singular and plural.

⁴Typically a temperature profile.

the observations, usually denoted y in an EnKF or 4d-VAR DAS context, while at the same time y in Rodgers (2000) denotes the channel radiances. Moreover in what follows we propose that we use a transformation of x that involves the AK (A) as an observation in the DAS, so to bridge the notation gap, we will denote this transformation as y_A .

4. Averaging kernel method

In theory (*e.g.*, Rodgers 2000), a significant fraction of retrieval error results from the fact that the retrievals can not observe the detailed vertical structure of the atmosphere and prior information is used to fill some of the gaps, but the result is smoothed compared to reality. When used as a profile, this corresponds to correlated errors. Also, using the best possible prior (the forecast!) is considered a no-no by many. Rodgers (2000, Section 8.3.4) suggests a one-sentence way around this. He leaves a few details for the student to fill in, but most of what we need is scattered throughout his excellent book. To begin, Rodgers (2000, Eq. 3.12) writes the retrieval as

$$\hat{x} = Ax + (I - A)x_a + G_y \varepsilon_y. \quad (1)$$

Here \hat{x} is the retrieval, x is the true state vector, and x_a is the prior. The matrix A is the averaging kernel. The total measurement error relative to the forward problem, ε_y (Rodgers 2000, Eq. 3.11), includes measurement error and forward model error and is usually characterized as being unbiased and having a specified covariance matrix, S_ε .⁵ In Eq. (1) the matrix G_y is the sensitivity of the retrieval to the radiances and is given by Rodgers (2000, Eqs. 3.25, 3.27) as

$$G_y = \frac{\partial \hat{x}}{\partial y} = \hat{S} K^T S_\varepsilon^{-1}. \quad (2)$$

Here $K = \frac{\partial F}{\partial x}$ is the Jacobian of the forward problem, F , evaluated at the solution \hat{x} (Rodgers 2000, Eq. 2.20), and

$$\hat{S}^{-1} = K^T S_\varepsilon^{-1} K + S_a^{-1} \quad (3)$$

defines the accuracy (inverse of the error covariance) of the retrieval in terms of the accuracy of $G_y y$ and the accuracy of x_a (Rodgers 2000, Eqs. 2.27, 4.13). An alternative computational formula for G_y , given by Rodgers (2000, Eqs. 2.45) is

$$G_y = S_a K^T S_q^{-1} \quad (4)$$

with

$$S_q = K S_a K^T + S_\varepsilon \quad (5)$$

in which S_q should be well conditioned due to the presence of S_ε .⁶ With this notation we may write

$$A = G_y K = \hat{S} (K^T S_\varepsilon^{-1} K) = I - \hat{S} S_a^{-1} \quad (6)$$

⁵In practice S_ε is often taken to be diagonal, but that is not material in this analysis. We return to a discussion of S_ε in § 6.

⁶Rodgers (2000, Section 4.1) demonstrates the algebraic equivalence of these two forms. In brief, find that $\hat{S}^{-1} G_y S_q = K^T (I + S_\varepsilon^{-1} K S_a K^T)$ for both forms of G_y .

following Rodgers (2000, Eqs. 2.78, 2.79).⁷

Within a generalized DAS, where any observation can be used if it can be calculated, this formulation of the retrieval can be used to eliminate the impact of the prior. For this purpose we posit a pre-processing step that removes the prior and effectively removes smoothing as an error source (Rodgers 2000, Section 8.3.4) by defining the new observation as

$$\hat{y}_A = \hat{x} - (I - A)x_a, \quad (7)$$

and the new observation operator by

$$y_A = Ax. \quad (8)$$

The retrievals are of course still smooth, but we will now be comparing observed and simulated quantities with the same degree of smoothing. Now the observation increments

$$\hat{y}_A - y_A = G_y \epsilon_y, \quad (9)$$

are unbiased if the ϵ_y are unbiased, and have covariance

$$S_m = G_y S_\epsilon G_y^T. \quad (10)$$

Rodgers (2000, Eq. 3.19) calls this the retrieval noise covariance. An alternative form for Eq. (10) is

$$S_m = A \hat{S}, \quad (11)$$

which follows from substituting Eq. (2) into Eq. (10), and then making use of the second form of Eq. (6).⁸

The end result is the same situation that is always obtained in maximum *a posteriori* (MAP) retrieval schemes—you compare an observation ($\hat{x} - (I - A)x_a$) to a simulated value (Ax) such that the mean difference has zero expectation and you normalize the difference by the square root of its inverse covariance ($S_m^{-\frac{1}{2}}$).

In this approach, the H -operator for this observation is simply Ax , but remember that x is defined in the space of the forward radiative transfer model. Therefore the observation operator will in general require time interpolation followed by interpolation in latitude and longitude, followed by vertical interpolation to the vertical grid of the forward problem and finally the computation of y_A according to Eq. (8). Then, each element of the observation vector, $\hat{y}_A(i)$, will be compared to the i th element of the simulated vector $y_A(i)$. For convenience we might store the elements of the observation vector individually, and compute the elements of the simulated vector individually as the dot product of the i th row of A and x . Practically then, the individual elements of the observation vector, $\hat{y}_A(i)$, might be stored along with the corresponding rows of A . The rows of A might also be used in determining vertical localization. Note that the preceeding arguments formally depend on the correctness of Eq. (1), including the implicit linearization assumption.

⁷To see the equivalence of the second two forms, note that $I = \hat{S}\hat{S}^{-1} = \hat{S}(K^T S_\epsilon^{-1} K + S_a^{-1})$ according to Eq. (3), and subtract the second form to find that $I - A = \hat{S}S_a^{-1}$.

⁸Although A is not symmetric, $A\hat{S}$ is symmetric since S_m is symmetric. Explicitly $(A\hat{S})^T = \hat{S}^T A^T = \hat{S}(K^T S_\epsilon^{-1} K)^T \hat{S} = \hat{S}(K^T S_\epsilon^{-1} K)\hat{S} = A\hat{S}$.

If we scale by $S_m^{-\frac{1}{2}}$ we effectively rotate to a space where the observation errors are unbiased, uncorrelated, and have unit variance. Specifically we redefine the new observation as

$$\hat{y}_A = S_m^{-\frac{1}{2}}(\hat{x} - (I - A)x_a), \quad (12)$$

and the new observation operator by

$$y_A = S_m^{-\frac{1}{2}}Ax = A_Rx, \quad (13)$$

where we have defined the “rotated” AK,

$$A_R = S_m^{-\frac{1}{2}}A, \quad (14)$$

for convenience. Now the observation increments

$$\hat{y}_A - y_A = S_m^{-\frac{1}{2}}G_y\epsilon_y, \quad (15)$$

are unbiased, and have covariance I .⁹

5. Forecast as prior

After the pre-processing of Eq. (12), x_a is no longer needed. Since the observation increment does not depend on the prior, we can use any prior. One possibility is to use the ensemble background mean as x_a and the ensemble sample covariance as S_a . We would call this the interactive retrieval approach. This complicates the data flow since now the background ensemble interpolated to the observation time, location, and vertical structure must be provided to the retrieval scheme. However, this is the best possible prior for the retrieval within the context of data assimilation. Further, if the system has settled down, then the ensemble background mean should be a good estimate of the truth and the linearization assumption of Eq. (1) should be a better approximation than otherwise.

6. Forward problem errors

The correct specification of the statistics of the combined measurement and forward problem errors, S_ϵ , is crucial to the method. Usually we think of S_ϵ as something specified by the retrieval-performing organization. However, there are two different approaches to determining S_ϵ , and the second one involves the data assimilation system. The first, what might be called a bottom-up approach, tries to estimate all the sources of error separately and build up a composite estimate of S_ϵ . The second, or top-down approach, is based on comparisons of radiances and simulated radiances, most conveniently within the context of a radiance data assimilation system.

In the bottom-up approach the following sources of error should be considered:

⁹Since S_m is symmetric, we have $\langle(\hat{y}_A - y_A)(\hat{y}_A - y_A)^T\rangle = S_m^{-\frac{1}{2}}G_y\langle\epsilon_y\epsilon_y^T\rangle G_y^T S_m^{-\frac{1}{2}T} = S_m^{-\frac{1}{2}}S_m S_m^{-\frac{1}{2}T} = I$.

- Measurement error including calibration error, misrepresentation of the instrument response function, sources of stray radiation, geometry error, geolocation error, polarization characterization, ...;
- Spectroscopy errors (in line widths, strengths, continuum, ...);
- Neglect of NLTE effects;
- Neglect or errors in how the magnetic field interacts with radiation, including Zeeman line-splitting, Faraday rotation, ...;
- Integration errors, i.e., truncation errors from the vertical grid;
- EOF truncation errors;
- Cloud clearing errors and beam-filling effects; and
- Errors in physical properties that must be specified, but are not part of the control vector.

The last item might include aerosol, cloud, and surface properties as well as trace gas absorber amounts, such as CO₂ amounts, and the cosmic background in the microwave. All of these effects could be estimated in simulation. Estimating cloud clearing errors is likely the most difficult.

Among the top-down approaches, those derived from the method of Desroziers et al. (2005) seem most promising. Desroziers et al. (2005) combine differences among the observations, the analysis, and the background, all in observation space to estimate the analysis, background, and observation error covariances—again, all in observation space. To use this approach to estimate the forward problem error would require the radiance observations and the background and analysis evaluated as radiances. For discussion here, let x_b be the background for one of the ensemble members interpolated to the observation time and location, and to the vertical grid of the forward problem.¹⁰ Then we can estimate the corresponding background radiance y_b as

$$y_b = \hat{y} + K(x_b - \hat{x}), \quad (16)$$

where $\hat{y} = F(\hat{x})$ is the radiance associated with the retrieval. The ensemble mean of y_b would be the appropriate background in observation space for the Desroziers et al. (2005) method. The same approach applies to the analysis radiance, except to note that the analysis is normally only available at the analysis time. However, in the LETKF approach, the analysis weights can be applied at each time (say each hour) to define the analysis at the same times as the background. Or the analysis weights could be applied directly to the y_b to define the analysis in observation space for the Desroziers et al. (2005) method.

Examples of studies that apply the so-called “Desroziers diagnostic” to radiance observation errors are Garand et al. (2007); Bormann and Bauer (2010); Bormann et al. (2010).¹¹ Within an

¹⁰This interpolation is part of our observation operator.

¹¹The later two references also apply two other methods, the so-called Hollingsworth/Lonnberg method (Hollingsworth and Lonnberg 1986) and a method based on subtracting a scaled version of mapped assumed background errors from FG-departure covariances.

ensemble Kalman filter, Li et al. (2009) demonstrate how it is possible to continuously update both the observation and background error statistics based on Desroziers et al. (2005).

For radiance data assimilation, the Desroziers et al. (2005) approach includes representative-ness errors in the radiance errors. If we don't want to include this error component, and Rodgers (2000) would not, then we can project the observed radiances onto the subspace sampled by the model. To do this we would need to assume that the forward model is correct to transfer model states to radiances, but then we would only use the leading EOFs of a diverse sample of model states in radiance space to filter the observations.

7. EOF properties

Definition. Empirical orthogonal functions (EOFs) can be defined in terms of any state vector. For this description, suppose the state vector x is the list of temperatures at the levels used in the forward problem. But x could also include other profile parameters and surface parameters with no loss of generality. The EOFs are the eigenvectors of the x sample-covariance matrix. The EOFs are ordered by eigenvalue. The variance explained by each EOF is given by its eigenvalue. Thus, the first EOF is the anomaly “pattern” that occurs most frequently.

Data compression. Since an EOF representation is often used in the retrieval to filter small scales, the retrieval solution can be represented equivalently (and with no loss of information) as a vector of EOF coefficient. In the case of our AER retrievals we use a single set of EOFs globally to avoid generating artificial edges in maps of the retrieved fields. For Mars TES, 4-6 EOFs are sufficient. Then the retrievals, error covariances, and AKs can be provided to the DAS in terms of the leading EOFs. If we reduce the size of the data vector from n (say 30) to j (say 5) we also reduce the size of the AK and error covariance matrices from $O(n^2)$ to $O(j^2)$.

EOF algebra. EOFs are used to relate anomalies, denoted x' , from an overall mean, \bar{x} . That is,

$$x' = x - \bar{x}. \quad (17)$$

In any implementation we first subtract \bar{x} before converting to EOF coefficients and then add it back when converting to physical space. The EOFs are truncated and ordered as columns in a matrix E . In what follows it is only important that $E^T E = I$. But note that if all EOFs are retained then E^T and E are inverses so that $EE^T = I$ as well.

The projection of x on the EOFs gives the EOF coefficients α according to

$$E^T x' = \alpha. \quad (18)$$

Reconstruction and filtering (in the case where we have truncated the EOF series) is obtained by

$$E\alpha = x'_f. \quad (19)$$

If we left multiply Eq. (19) by E^T , then, since $E^T E = I$, we find $E^T x'_f = \alpha$. Therefore $E^T (x' - x'_f) = 0$, demonstrating that the part of x' that is filtered does not project onto the retained EOFs.

In summary, we made the following definitions:

- x' is the state vector [K]. (Note that this is a deviation from an overall mean.)
- E is the matrix of EOFs [1].
- α is the vector of EOF coefficients [K].

We established these relationships:

- $E^T E = I$ is the orthonormality relationship [1].
- $E^T x' = \alpha$ is the projection method [K].
- $E \alpha = x'$ is the filter or reconstruction method [K].

EOF covariances. The sample covariance matrix of x is

$$S = \langle (x - \langle x \rangle)(x - \langle x \rangle)^T \rangle = E \langle (\alpha - \langle \alpha \rangle)(\alpha - \langle \alpha \rangle)^T \rangle E^T = E \tilde{S} E^T \quad (20)$$

since $x - \langle x \rangle = (x - \bar{x}) - (\langle x \rangle - \bar{x}) = E\alpha - E\langle \alpha \rangle$. Here angle brackets denote an average over the sample. If the sample is the posteriori distribution then $\langle x \rangle = \hat{x}$, if the sample is the prior distribution then $\langle x \rangle = x_a$, and if the sample is the climate distribution then $\langle x \rangle = \bar{x}$ and $\langle \alpha \rangle = 0$. If E includes all eigenvectors, then $\tilde{S} = E^T S E$. Generally we truncate the EOF series so E is not full rank and as a result, if the retrieval is done in terms of EOF coefficients, then \hat{S} will not be full rank either. However, we can still apply $\tilde{S} = E^T S E$ in this case to obtain the upper left corner of the full \tilde{S} matrix as desired.¹² While Eq. (20) can be used to transform from EOF to physical space covariance it is better to add in the “noise” present in the truncated EOFs. For example, in the AER retrieval we determine the error covariance matrix of the EOF coefficients.¹³ Since the coefficients of the truncated EOFs are not estimated during the retrieval we can add back the variability of these EOFs given by the associated eigenvalues. Computationally, take the diagonal matrix of eigenvalues, replace the upper left block with the estimated error covariance of the EOF coefficients from the retrieval scheme and apply Eq. (20) using the full (not truncated) matrix E .

8. EOF analysis of the AK method

Note that Eq. (1) is usually written for the full physical space quantities, but we can subtract $\bar{x} = A\bar{x} + (I - A)\bar{x}$ from Eq. (1) to get a version in terms of anomalies,

$$\hat{x}' = Ax' + (I - A)x'_a + G_y \epsilon_y. \quad (21)$$

We now project Eq. (21) into EOF coefficient space. To do this we filter each x' -vector. For justification, note that the filtering step is normally done within the retrieval algorithm initially and at each step of the iteration. Since putting too much structure in the background for the retrieval

¹²Partition E into retained and truncated EOFs $[E_r \ E_t]$ to see this.

¹³This is a full symmetric matrix.

can degrade the results, it is also helpful to filter the prior. Further, within the context of a retrieval in terms of EOF coefficients, that part of the true x' that does not project on the EOFs should be considered part of the representativeness error.

To filter each x' -vector we simply replace each x' by its filtered version $EE^T x'$. Then, to project into the EOF space, we left multiply each equation by E^T . In the case of Eq. (21) we obtain

$$(E^T E)(E^T \hat{x}') = (E^T A E)(E^T x') + [E^T (I - A)E](E^T x'_a) + E^T G_y \varepsilon_y, \quad (22)$$

where we have collected terms in parentheses to indicate how to apply the above relationships. If we define

$$\tilde{A} = E^T A E \quad (23)$$

then we have

$$\hat{\alpha} = \tilde{A} \alpha + (I - \tilde{A}) \alpha_a + \varepsilon_\alpha, \quad (24)$$

where

$$\varepsilon_\alpha = E^T G_y \varepsilon_y, \quad (25)$$

$$\hat{\alpha} = E^T (\hat{x} - \bar{x}), \quad (26)$$

and

$$\alpha_a = E^T (x_a - \bar{x}). \quad (27)$$

We may say that \tilde{A} is the projection of A into the EOF coefficient space, or simply that \tilde{A} is the EOF AK, and that ε_α is the projection of the retrieval error into the EOF coefficient space, or simply that ε_α is the EOF retrieval error.

Now ε_α has covariance given by

$$\tilde{S}_m = E^T G_y \langle \varepsilon_y \varepsilon_y^T \rangle G_y^T E = E^T G_y S_\varepsilon G_y^T E = E^T S_m E. \quad (28)$$

As we did before (in § 4), we scale by $\tilde{S}_m^{-\frac{1}{2}}$ to rotate to a space where the observation errors are unbiased, uncorrelated, and have unit variance. Specifically we redefine the new observation as

$$\hat{y}_A = \tilde{S}_m^{-\frac{1}{2}} (\hat{\alpha} - (I - \tilde{A}) \alpha_a), \quad (29)$$

and the new observation operator by

$$y_A = \tilde{S}_m^{-\frac{1}{2}} \tilde{A} \alpha = \tilde{S}_m^{-\frac{1}{2}} \tilde{A} E^T x' = \tilde{A}_R (x - \bar{x}), \quad (30)$$

where we have defined the “rotated” EOF-space AK,

$$\tilde{A}_R = \tilde{S}_m^{-\frac{1}{2}} \tilde{A} E^T, \quad (31)$$

for convenience. Now the observation increments

$$\hat{y}_A - y_A = \tilde{S}_m^{-\frac{1}{2}} \varepsilon_\alpha, \quad (32)$$

are unbiased, and have covariance I .

Note that we include the EOF projection operator in our definition of \tilde{A}_R in Eq. (31). This allows for the possibility of vertical localization and makes the implementation in EOF or physical space similar. Alternatively, one could determine α by projection (as in Eq. (26) and Eq. (27)), store rows of $\tilde{S}_m^{-\frac{1}{2}}\tilde{A}$ with the observation \hat{y}_A and then define the observation operator by first form in Eq. (30). This approach saves storage since it stores weights for EOF coefficients, not temperatures, but is not amenable to localization. Of course vertical localization may not be possible depending on the structure of \tilde{A}_R .

9. Using the EOF representation for vertical interpolation

Vertical interpolation inevitably introduces some errors. If we can somehow project the meteorological model vertical structure directly onto the EOFs we could eliminate this source of error. In this case we would use the first form of Eq. (30), $\tilde{S}_m^{-\frac{1}{2}}\tilde{A}\alpha$. As described above we must interpolate the meteorological model values to the levels used in the retrieval and then project using Eq. (18). However there are some complications that make this unworkable when the grid definitions are very different, *e.g.*, altitude relative to the geoid vs. sigma. What follows is an alternative approach (that, as a by-product, provides a method of vertical interpolation).

In general we want to convert a vertical profile of model temperatures into the EOF coefficients defined relative to the retrieval vertical coordinate. For example, in the model, temperature might be stored as σ -layer values, while the retrieval might operate on a fixed p -level grid. In this example, at some locations some of the p -levels might be below model topography. We require (and this is expected to always be the case) that there are more model temperatures in the profile than there are EOF coefficients. To convert from model temperatures, T , to EOF coefficients, α , we determine α to fit the T in a least squares sense. That is we minimize

$$J = \sum_i (T_i^* - T_i)^2, \quad (33)$$

where T_i are the background model temperatures and T_i^* are the model temperatures reconstructed from the α . Here we restrict the T_i to just those model temperatures that can be reconstructed from a specification of the EOF coefficients. For example, any model temperatures above the uppermost temperature in the retrieval vertical coordinate would not be included in the sum in Eq. (33). The reconstructed temperatures are determined from α in two steps. (We don't actually perform these steps, but we must define them to solve the least squares problem.) First we reconstruct the temperature profile x in the retrieval vertical coordinate,

$$x = E\alpha + \bar{x}. \quad (34)$$

Second we interpolate to the model vertical coordinate using an interpolation operator V that is linear in the x ,

$$T^* = Vx = V(E\alpha + \bar{x}). \quad (35)$$

Usually we will interpolate the temperatures linearly in pressure. Note that in the example described above, V will depend on the model surface pressure and vary from location to location. Instead of doing an actual interpolation, we must be able to determine the matrix V . To minimize J with respect to the α we set the derivatives equal to zero and find that¹⁴

$$E^T V^T V E \alpha = E^T V^T (T - V \bar{x}). \quad (36)$$

This is a linear equation of the form $M\alpha = b$ where the size of M might be 6×6 , but M will vary from location to location. Note in Eq. (36) that it is only the (presumably smooth) \bar{x} that is vertically interpolated.

Given the α determined this way we can then calculate x according to Eq. (34).¹⁵ This would fit into either the AK method or the EOF-AK method.

10. Implementation details

a. Underground when using pressure as the vertical coordinate

If the radiative transfer model (RTM) for F uses a grid fixed in pressure (or in altitude) then the lengths of the vectors and the shapes of the matrices will vary from retrieval to retrieval as topography and surface pressure vary from location to location. We may simply truncate those parts of the vectors and matrices that are underground. Operationally, we may define underground as those levels whose corresponding rows of $K \equiv 0$, *i.e.*, levels that have no effect on the calculated radiances. Alternatively, we can keep all the vectors and matrices full size but take care that underground levels are handled properly. For example, underground rows of K must be zero or must be strictly ignored. In particular, underground elements of S_m , $S_m^{-\frac{1}{2}}$, and S_a^{-1} should be reset to zero after they are calculated to avoid the possible effects of accumulating round-off errors.

When using EOFs, underground levels in A , G_y or S_m , and each x' should be zero. Equivalently the variables in Eq. (24) will be assured to be correct by setting the rows of E below the surface to zero.

b. Generalized matrix inversion

Since S_a may be poorly conditioned it will generally be necessary to use a truncated eigenvector decomposition to find the (Moore-Penrose) inverse of S_a . This is not necessary if \hat{S} is provided along with \hat{x} . Alternatively, if we are given S_ϵ , S_a , and K , it is possible to sidestep \hat{S} entirely by using Eq. (4) and Eq. (5). Similarly to compute $S_m^{-\frac{1}{2}}$ from S_m it may be necessary to use a truncated eigenvector decomposition.

¹⁴Using Einstein notation (summation implied for repeated indices), the reconstructed temperature at level i is written as $T_i^* = V_{ik}(E_{kn}\alpha_n + \bar{x}_k)$. The derivative of this with respect to α_j may be written as $E_{jk}^T V_{ki}^T$. Then setting the derivative of J to zero gives $0 = 2E_{jk}^T V_{ki}^T (T_i^* - T_i)$. Substituting for T^* here and rearranging terms then gives Eq. (36).

¹⁵Note that this produces temperature values for all radiative transfer model grid points, even below the surface in some cases. Any values below the surface will not be used by the radiative transfer model, but by construction these will be statistically consistent with the rest of the profile.

The idea here is that after we rotate to the eigenvector coordinate system there are some coordinates that are so flat (have such small eigenvalues) that they can be ignored. So we determine an inverse or square root inverse that is formally correct in the full but transformed space and then ignore (truncate) the near-singular dimensions. The inverse matrix in the truncated transformed coordinate system is in fact a proper Moore-Penrose pseudoinverse matrix in the original coordinate system. The correct choice of truncation would be the truncation used in the retrieval scheme. If an EOF truncation is not used in the retrieval scheme then the correct choice of truncation may be apparent from the eigenvalue spectrum. Otherwise it will probably be sufficient to choose the truncation such that at least 99.9% of the variance is retained, *i.e.*, such that the sum of the retained eigenvalues is $\geq 99.9\%$ of the sum of all the eigenvalues.

Eigen-decomposition. Any positive definite matrix S , such as a covariance matrix, may be written as

$$S = VGV^T, \quad (37)$$

where V is a matrix of the eigenvectors and G is a diagonal matrix of the eigenvalues of S .

Inverse. The pseudoinverse of S is given by

$$S^+ = VG^{-1}V^T \quad (38)$$

since

$$SS^+ = VGV^T VG^{-1}V^T = VV^T = I. \quad (39)$$

Note that $V^T V = I$ regardless of the truncation, but $VV^T = I$ is true only for the case of no truncation. When we truncate both V and G we can easily demonstrate that S^+ satisfies the four required Moore-Penrose properties— $SS^+S = S$, $S^+SS^+ = S^+$, $(SS^+)^T = SS^+$, and $(S^+S)^T = S^+S$.¹⁶

Square root. The square root of S is given by

$$S^{\frac{1}{2}} = VG^{\frac{1}{2}}V^T \quad (40)$$

since

$$S^{\frac{1}{2}}S^{\frac{1}{2}} = VG^{\frac{1}{2}}V^T VG^{\frac{1}{2}}V^T = VGV^T = S \quad (41)$$

for any truncation. According to Eq. (38) the pseudoinverse of $S^{\frac{1}{2}}$ is then

$$S^{\frac{1}{2}+} = VG^{-\frac{1}{2}}V^T. \quad (42)$$

c. Interface with retrieval and order of computation

Consider how data flows through our procedures for a moment. Given the retrieval, \hat{x} , once we know A and S_m , we then calculate \hat{y}_A , and A_R from Eq. (12), and Eq. (14). Usually, the retrieval will not provide A and S_m . Three possible pathways for the calculation of A and S_m are described

¹⁶From Eq. (39) we have that $SS^+ = VV^T$ which is symmetric, and reversing the roles of S and S^+ in Eq. (39) gives the same result. Thus SS^+ and S^+S are symmetric and equal to VV^T . Clearly pre- or post-multiplying S or S^+ by VV^T leaves them unchanged, demonstrating the first and second properties. And because of symmetry the third and fourth properties hold.

| Pathway | 1 | 2 | 3 |
|---------------------|-----------------------------------|-------------------------------------|-----------------------|
| Interface | | | |
| Retrieval inputs | x_a, S_a, S_ε | x_a, S_a | x_a |
| Retrieval outputs | \hat{x}, K | \hat{x}, \hat{S} | \hat{x}, \hat{S}, A |
| Calculations | | | |
| $S_q =$ | $KS_aK^T + S_\varepsilon$ [Eq. 5] | — | — |
| $G_y =$ | $S_aK^TS_q^{-1}$ [Eq. 4] | — | — |
| $A =$ | G_yK [Eq. 6] | $I - \hat{S}\hat{S}_a^{-1}$ [Eq. 6] | — |
| $S_m =$ | $G_yS_\varepsilon G_y^T$ [Eq. 10] | $A\hat{S}$ [Eq. 11] | $A\hat{S}$ [Eq. 11] |

Table 1: Pathways to calculate A and S_m . See text for discussion.

now. (Refer to Table 1.) Inputs to the retrieval include x_a , S_a , and S_ε . In some situations, these three inputs are all constant. The three sets of outputs from the retrieval are \hat{x} and (i) K , or (ii) \hat{S} , or (iii) \hat{S} and A . In the first pathway, from the retrieval inputs S_a and S_ε and the specified retrieval output K we can calculate in turn S_q , G_y , A , and S_m from Eq. (5), Eq. (4), the first form of Eq. (6), and Eq. (10). This pathway requires only the inversion of S_q , which should be well conditioned as mentioned before. In the second pathway, using the retrieval input S_a , and the retrieval output \hat{S} we can calculate A from the third form of Eq. (6) and then S_m from Eq. (11). In this case we need to invert S_a . As pointed out by Deeter et al. (2003) in cases where x_a and S_a are fixed, this second pathway greatly reduces the data that must be provided by the retrieval and S_a^{-1} needs to be calculated only once. In the third pathway using the retrieval outputs \hat{S} and A we can immediately calculate S_m from Eq. (11). Note that in the second and third pathways the retrieval outputs are all sized by the number of retrieved quantities, and not the number of radiances.

When using EOFs, the above discussion still holds, but now we also need \bar{x} and E to define the EOF representation. Then once we have determined A and S_m as before, we can then calculate in turn \tilde{A} , $\hat{\alpha}$, α_a , \tilde{S}_m , \hat{y}_A , and \tilde{A}_R from Eq. (23), Eq. (26), Eq. (27), Eq. (28), Eq. (29), and Eq. (31).

11. Algorithm implementation

Here I sketch out *one* implementation for interactive retrievals using the AK-EOF method and the EOF vertical interpolation operator.

1. Assemble known constants: instrument description (channels, radiance errors (S_ε ¹⁷), geometry), EOF information (EOFs (E), climate mean (\bar{x}), truncation, eigenvalues),
2. Build observation file with times, locations, radiances (y), geometry,

¹⁷This should be instrument noise plus forward problem errors (Rodgers 2000, Eq. 3.11). However in practice just instrument noise is often used.

3. Use the background ensemble to calculate the prior mean and covariance, x_a and S_a , and add these to the observation file. Note that x_a and \bar{x} are different in this case.¹⁸
4. Perform retrieval. Add \hat{x} and \hat{S} in the retrieval output file.¹⁹
5. Post-process retrieval. Following pathway 2, calculate in turn A , S_m , \tilde{S}_m , \tilde{A} , $\hat{\alpha}$, α_a , \hat{y}_A (from Eq. (29)), and \tilde{A}_R . Save \hat{y}_A and \tilde{A}_R in (LETKF) obs-data structure. The i th element of \hat{y}_A is one observation. It is associated with the i th row of \tilde{A}_R .
6. Data selection. For vertical data localization use the absolute value of the i th row of A_R as weights in the same way that we now localize radiances.
7. Vertical interpolation. After interpolating the model grids to the map location and time of the observation, use the procedure of § 9 to determine α and then x using Eq. (36) and Eq. (34).
8. Simulate observation. Calculate y_A according to Eq. (30).²⁰

The case outlined above is the most complicated one. Other implementations of the methods described here could follow this plan with appropriate simplifications.

12. Acknowledgments

I thank Dan Gombos, Janusz Eluskiewicz, Steven Greybush, Matthew J. Hoffman, Kayo Ide, Alan Lipton, Eugenia Kalnay, Jean-Luc Moncet, David Kuhl, Thomas Nehrkorn, Vivienne Payne, and R. John Wilson for helpful discussions. This work supported in part by NASA grant NNX07AN97G.

References

Bormann, N. and P. Bauer, 2010: Estimates of spatial and interchannel observation-error characteristics for current sounder radiances for numerical weather prediction. I: Methods and application to ATOVS data. *Quart. J. Roy. Meteor. Soc.*, **136** (649), 1036–1050.

¹⁸We retain the global EOFs in our process instead of trying to determine EOFs for each retrieval—we want the EOFs to span all possibilities. The alternative, a set of EOFs from a forecast, might be missing a mode of variability actually present in reality and such an inconsistency could adversely affect the retrievals. I believe this issue—the fact that the ensemble does not have enough variability to efficiently fit the radiances—needs some sort of fix in radiance assimilation. Quality control might help, but is the wrong solution—getting rid of good but unusual data because the model is inadequate.

¹⁹We would also save K and $\hat{y} = F(\hat{x})$ as defined in § 6 to be used later in Eq. (16) if we are applying the Desroziers et al. (2005) method.

²⁰At this point we could also calculate the background radiance, y_b , as described in § 6 and save it in the observation data structure to be used later to define the analysis and background “observations” for use in the Desroziers et al. (2005) method.

- Bormann, N., A. Collard, and P. Bauer, 2010: Estimates of spatial and interchannel observation-error characteristics for current sounder radiances for numerical weather prediction. II: Application to AIRS and IASI data. *Quart. J. Roy. Meteor. Soc.*, **136** (649), 1051–1063.
- Deeter, M. N., L. K. Emmons, G. L. Francis, D. P. Edwards, J. C. Gille, J. X. Warner, B. Khatatov, D. Ziskin, J.-F. Lamarque, S.-P. Ho, V. Yudin, J.-L. Attié, D. Packman, J. Chen, D. Mao, and J. R. Drummond, 2003: Operational carbon monoxide retrieval algorithm and selected results for the MOPITT instrument. *J. Geophys. Res.*, **108** (D14), 4339–4349, doi:10.1029/2002JD003186.
- Desroziers, G., L. Berre, B. Chapnik, and P. Poli, 2005: Diagnosis of observation, background and analysis-error statistics in observation space. *Quart. J. Roy. Meteor. Soc.*, **131**, 3385–3396, doi:10.1256/qj.05.108.
- Eyre, J. R., G. A. Kelly, A. P. McNally, E. Andersson, and A. Persson, 1993: Assimilation of TOVS radiance information through one-dimensional variational analysis. *Quart. J. Roy. Meteor. Soc.*, **119**, 1427–1463.
- Garand, L., S. Heilliette, and M. Buehner, 2007: Interchannel error correlation associated with AIRS radiance observations: Inference and impact in data assimilation. *J. Appl. Meteor. Climatol.*, **46**, 714–725, doi:10.1175/JAM2496.1.
- Hoffman, R. N., 1983: Three-dimensional inversion of satellite observed radiances: A proposal. *Fifth Conference on Atmospheric Radiation*, American Meteorological Society, Boston, MA, Baltimore, Maryland, 43–46.
- Hoffman, R. N. and T. Nehr Korn, 1989: A simulation test of three-dimensional retrievals. *Mon. Wea. Rev.*, **117**, 473–494.
- Hollingsworth, A. and P. Lönnberg, 1986: The statistical structure of short-range forecast errors as determined from radiosonde data. Part I: The wind field. *Tellus A*, **38** (2), 111–136.
- Joiner, J. and A. M. da Silva, 1998: Efficient methods to assimilate remotely sensed data based on information content. *Quart. J. Roy. Meteor. Soc.*, **124** (549), 1669–1694.
- Li, H., E. Kalnay, and T. Miyoshi, 2009: Simultaneous estimation of covariance inflation and observation errors within an ensemble Kalman filter. *Quart. J. Roy. Meteor. Soc.*, **135** (639), 523–533.
- Pavelin, E. G., S. J. English, and J. R. Eyre, 2008: The assimilation of cloud-affected infrared satellite radiances for numerical weather prediction. *Quart. J. Roy. Meteor. Soc.*, **134** (632), 737–749.
- Rodgers, C. D., 2000: *Inverse Methods for Atmospheric Sounding: Theory and Practice*, Atmospheric, Oceanic and Planetary Physics, Vol. 2. World Scientific, Singapore.

Szyndel, M. D. E., A. D. Collard, and J. R. Eyre, 2004: A simulation study of 1D variational cloud retrieval with infrared satellite data from multiple fields of view. *Quart. J. Roy. Meteor. Soc.*, **130**, 1489.